arXiv:1612.03709v1 [cs.PF] 12 Dec 2016

Geographical Load Balancing across Green Datacenters: a Mean Field Analysis*

Giovanni Neglia
Université Côte d'Azur - Inria, France
giovanni.neglia@inria.fr

Matteo Sereno
Dip. di Informatica,
Università di Torino, Italy
matteo.sereno@di.unito.it

Giuseppe Bianchi
Dip. di Ingegneria Elettronica,
Università di Roma "Tor Vergata", Italy
giuseppe.bianchi@uniroma2.it

June 2016

Abstract

"Geographic Load Balancing" is a strategy for reducing the energy cost of data centers spreading across different terrestrial locations. In this paper, we focus on load balancing among micro-datacenters powered by renewable energy sources. We model via a Markov Chain the problem of scheduling jobs by prioritizing datacenters where renewable energy is currently available. Not finding a convenient closed form solution for the resulting chain, we use mean field techniques to derive an asymptotic approximate model which instead is shown to have an extremely simple and intuitive steady state solution. After proving, using both theoretical and discrete event simulation results, that the system performance converges to the asymptotic model for an increasing number of datacenters, we exploit the simple closed form model's solution to investigate relationships and trade-offs among the various system parameters.

1 Introduction

Providers such as Amazon, Google, Facebook, etc., are making a considerable effort to offer efficient, scalable, and reliable services. To achieve these

*Published in the proceedings of Greenmetics workshop 2016, in conjunction with ACM Sigmetrics/IFIP Performance, Juan-les-Pins, France, June 14, 2016

goals such services need to be supported by massive datacenters and relevant infrastructures to distribute power and provide cooling. Power management is becoming a crucial issue. Not only power consumption is ever increasing with an increasing user base and service expansion, but, as pointed out by several studies, the power consumption of datacenters is largely wasted.

In this paper we consider a set of micro-datacenters which are additionally powered by renewable energy sources, e.g., photovoltaic (PV) panels. Due to the current high costs for storing energy, the best use of renewable energy is to consume it when it is produced. Hence, we would ideally wish to adapt each micro-datacenter's load to the instantaneous energy production. One way to address such goal is to *federate* several micro-datacenters each other, and use a central controller to dispatch jobs where renewable energy is available, so as to minimize the (non-renewable) energy cost. The possibility to manage more jobs obviously offers a higher flexibility. The law of large numbers guarantees indeed that the aggregated load will be more regular and then easier to exploit for smart load scheduling as it is the case in a big datacenter. But when local renewable sources are available, a micro-datacenters' federation offers an additional advantage in comparison to a large datacenter: renewable energy production at different locations can be loosely correlated and then the aggregated energy production exhibits less variability.

Consider the following ideal case: a set of N identical datacenters, each with independent job arrival processes with rate λ and a single server with computing rate μ, and PV panels able to feed the datacenter a fraction s of the time. Compare it with a single datacenter which aggregates locally all the jobs as well as the computing and energy production infrastructure. The total normalized load for the federation of datacenter is $\rho = (N\lambda)/(N\mu) = \lambda/\mu$ with a normalized variability (standard deviation of the number of working servers divided by the average number of working server) equal to $\sqrt{(1-\rho)/(\rho N)}$. Similarly the federation can power through renewables a fraction s of its computing resources with a normalized variability equal to $\sqrt{(1-s)/(sN)}$, if the amounts of renewable energy produced at different datacenters can be considered independent. The single datacenter manages the same aggregate load with the same normalized variability, but the situation is different energywise. The single datacenter can be powered by renewables a fraction s of the time, but now the normalized variability is $\sqrt{(1-s)/s}$, if, as it is reasonable to assume in first approximation, all the PV panels at a given location produce (/do not produce) at the same time.

The example above is clearly over-simplified, it ignores the costs of job dispatchment among the micro-datacenters, the effect of fixed energy costs that are easier to optimize at a single datacenter, the possibility that renewable energy dynamics are too fast to be exploited by smart scheduling strategies, how revenues should be split among the datacenters, etc. Nevertheless, this example highlights a potential benefit from federating micro-

datacenters, that is interesting to quantify. As we are going to show below, even simple models for job traffic and energy production lead soon to scenarios for which it is difficult to provide closed-form expressions for the energy cost of a federation of micro-datacenters. One may then need to rely on expensive simulations that hide the role played by the different parameters. For this reason, in this paper we propose a mean field (fluid) model that is asymptotically correct and allows us to derive simple formulas for the main performance metrics, like the expected energy costs of the system.

The paper is organized as follows. After a brief discussion of related work, we introduce the system model in section III, and we provide and justify with both theoretical and simulation results a mean field approximation in section IV. In section V, we exploit the resulting simple model to quantify performance and trade-offs emerging in scenarios characterized by variable renewable energy production across micro datacenters.

2 Related Work

In Geographic Load Balancing (GLB) systems user requests are initially accepted by front-end elements and then redirected by a scheduler to geographically distributed datacenters for processing. The scheduler's decisions may depend on several mutually interacting (and in some case conflicting) objectives such as minimizing the electricity cost, the carbon-footprint and the response time. The paper [12] is one of the first studies about GLB. In particular, it focusses the attention on the key issues fostering the use of GLB such as different energy markets (e.g., day-ahead and real-time markets), and temporal or geographical energy price variations. The GLB represents the combination of these basic ingredients with the use of energy related metrics in the scheduler decisions. In this manner it is possible to account for different workload conditions, and time and geographical variability of the electricity costs.

In the last years other studies addressed the same problem by adding different scheduling constraints and/or by optimizing different metrics (see [14], [15], [16], [6], [10], [11]). For instance, the papers [14] and [11] introduce additional constraints for accounting QoS guarantees; while the interaction between GLB and smart grids, and then the exploitation of the workload demand-response capability have been addressed in [14]. Furthermore, the interaction of energy storage systems and GLB has been addressed in [6]. Indeed, storage systems can be used to smooth the variability of power supply and this is very important when the datacenters are powered by renewable sources.

Several studies pointed out that large datacenters are extremely expensive to maintain and this has encouraged the development of architectures that interconnects multiple micro-datacenters [3]. This trend influenced our

work because workload scheduling among a large number of interconnected datacenters gives rise to computational problems (e.g., see the summary of the techniques used in geographical load balancing in [13]).

The works closest to ours are [11] and [10], where geographical load balancing is driven by time-varying energy prices, that can be due to a significant local production from renewable sources. While in these papers energy prices are considered to be known in advance over some future time-horizon, in our case renewable energy production is a stochastic process and scheduling is decided on the basis of the current state of the system.

3 Problem

We consider a federation of N identical micro datacenters. The aggregated job arrival process at the federation is modeled as a Poisson process with rate $N\lambda$. The service time of each job is assumed to be exponential with expected value $1/\mu$.[1] Each datacenter is connected to the grid but it can be powered also by some renewable source. We consider here that the renewable source can be in two states: in state $\mathcal{S}$ (sunny) the energy produced by the source is able to power the whole datacenter, in state $\mathcal{C}$ (cloudy) the energy produced is negligible. Renewable states evolve according to a continuous time Markov Chain. Let ν_C and ν_S denote respectively the transition rates from $\mathcal{S}$ to $\mathcal{C}$ and from $\mathcal{C}$ to $\mathcal{S}$. The model for the renewable source can be made arbitrarily more realistic by adding multiple states. For the moment we assume that the Markov chains associated to renewable sources at different datacenters evolve independently.

When a new job arrives the scheduler dispatches it i) to a datacenter that is available to process it and in state $\mathcal{S}$ (i.e., currently powered by renewables) if any, otherwise ii) to an available datacenter if any, and as last option iii) to a central waiting queue from which the job will be moved to the first available datacenter. The system is then operating as an $M/M/N$ queue with the characteristic that available servers in state $\mathcal{S}$ get jobs with strict higher priority than other servers. Among the work conserving disciplines this intuitively minimizes the total expected energy cost.

The system can be described as a continuous time Markov Chain with state $(J^N(t), S^N(t), B_S^N(t))$, where $J^N(t)$ is the number of jobs in the system, $S^N(t)$ is the number of servers in state $\mathcal{S}$, and $B_S^N(t)$ is the number of servers busy (i.e., serving a job) and in state $\mathcal{S}$, all at time t. The Markov chain has a very particular structure: for example $J^N(t)$ itself, representing the number of jobs in a $M/M/N$ queue, evolves as a Markov chain. $S^N(t)$ is described by a simple Markov chain too. In particular the stationary dis-

[1] While we need an underlying Markovian process to correctly derive the asymptotic fluid model, empirical results show that the fluid model does not heavily depend on many of these assumptions.

tributions of $J^N(t)$ and $S^N(t)$ can be derived easily in closed-form. Despite these properties, it is not easy to characterize $B_S^N(t)$ and in particular we have not been able to derive in closed-form its stationary distribution. This is less surprising if we think about a similar problem for parallel queues where the simple join-the-shortest-queue policy couples the status of the different queues so that their stationary distribution can be expressed only as an infinite mixture of geometric distributions [1] (there are many works on priority queues and/or shortest queue policies, see for instance [4] [7]). Similarly, here our dispatching policy couples the two different states of a server (being busy and being powered by renewables) in a non-trivial way so that it is difficult to characterize the process $B_S^N(t)$, as we need to quantify the energetic savings coming from the federation.

In order to study the system $(J^N(t), S^N(t), B_S^N(t))$ we could resort to simulations or to a numerical solution of the Markov Chain. In both cases the computational cost increases with the number of datacenters N. These difficulties are aggravated if more realistic and then more complex models for traffic arrival process or renewable energy evolution are considered with a potential explosion of the state space. Moreover, the effect of the different parameters can be more difficult to unveil using numerical methods. For these reasons, as it has been successfully done in other fields, we derive the fluid limit of the Markov chain of interest, that allow us to obtain simple closed-form expressions for the main performance metrics independently from the system size N.

4 Fluid Model

In this section we show that the stochastic dynamics of the Markov chains $(J^N(t), S^N(t), B_S^N(t))$ converge in probability to a deterministic process as N diverges.[2] More precisely, we will show that if $1/N(J^N(0), S^N(0), B_S^N(0))$ converges to the constant values $(j_0, s_0, b_{s,0})$ when N diverges, then there exists a vector of deterministic functions $(j(t), s(t), b_s(t))$ such that $(j(0), s(0), b_s(0)) = (j_0, s_0, b_{s,0})$ and for any $T > 0$:

$$\sup_{0 \le t \le T} \left|\left| \frac{1}{N}(J^N(t), S^N(t), B_S^N(t)) - (j(t), s(t), b_s(t)) \right|\right| \underset{P}{\to} 0,$$

i.e., the rescaled process converges to $(j(t), s(t), b_s(t))$.

This kind of convergence results has become popular since the seminal work of Kurtz (see for example [9]), that shows that the limiting process can be described by a system of differential equations: $d\mathbf{x}/dt = \mathbf{f}(\mathbf{x}(t))$, where $\mathbf{f}()$ is called the limiting drift function. Classic results require $\mathbf{f}()$ to be a Lipschitz function. By carrying out the usual derivation of the fluid limit

[2] In what follows, convergence of random variables is always "in probability." We omit to repeat it at each time.

for the process $(J^N(t), S^N(t), B_S^N(t))$, the corresponding function $\mathbf{f}()$ will appear to be discontinuous and then it has not the Lipschitz property. Nevertheless, we can apply more recent and general results from [5] to show that the dynamics converge to the solution of a system of differential inclusions, i.e. where the function $\mathbf{f}()$ is replaced by a set valued function.

As we observed in the previous section, the processes $J^N(t)$ and $S^N(t)$ are themselves Markov chains. Rather than studying the joint system $(J^N(t), S^N(t), B_S^N(t))$ we first derive the fluid limits for $J^N(t)$ and $S^N(t)$ and then move to consider the fluid limit for $B_S^N(t)$. While we could directly consider the limit of the triplet, this approach can result easier to follow for the reader unfamiliar with fluid limits. Moreover, the results for $J^N(t)$ and $S^N(t)$ do not require the more complex machinery of differential inclusions, so this approach allows us to better highlight where difficulties arise for $B_S^N(t)$.

The Markov chain describing $J^N(t)$ is such that the transition from state J to state $J+1$ occurs with rate $N\lambda$, while the transition from state J to state $J-1$ occurs with rate μJ, if $J \leq N$, and with rate μN, if $J > N$. We consider now the scaled process $J^N(t)/N$, whose transition rates from x to $x + l/n$ can be expressed as $N\beta_l(x)$ where the functions $\beta_l(x)$ do not depend on N. In particular $\beta_1(x) = \lambda$, $\beta_{-1}(x) = \mu x$ for $x \leq 1$, $\beta_{-1}(x) = \mu$ for $x > 1$ and $\beta_l(x) = 0$ otherwise. The rate of changes of $J^N(t)/N$ is then

$$f(x) = \frac{1}{N}\left(\beta_1(x) - \beta_{-1}(x)\right) = \begin{cases} \lambda - \mu x & \text{if } x \leq 1 \\ \lambda - \mu & \text{if } x > 1 \end{cases}$$

that is a Lipschitz function. This property and the fact that $\sum_l |l|\beta_l(x) < \infty$ guarantee [9] that if $J^N(0)/N$ converges to j_0, $J^N(t)/N$ converges to the unique solution[3] of the following equation

$$\frac{dj}{dt} = f(j(t)), \quad j(0) = j_0. \tag{1}$$

Observe that $j(t) > 1$ corresponds to $J^N(t) > N$ and then a situation where all the N data centers are working and there are $J^N(t) - N$ jobs in the queue. Given that $f(x) < 0$ for $x > 1$, Eq. (1) shows that $j(t) < 1$ for large enough t and then after some transient the job queue is asymptotically empty and the number of jobs in the system coincides with the number of busy servers. Moreover, $j(t)$ converges when t diverges: $j^* \triangleq j(\infty) = \lambda/\mu = \rho$. This value is the only accumulation point for the possible trajectories of $j(t)$ and then it is also the stationary probability that a server is busy in the original Markov chain [2] (as it is known from the analysis of the $M/M/N$ queue).

[3] Continuity of the right hand side guarantees the existence of the solution and Lipschitz property guarantees uniqueness.

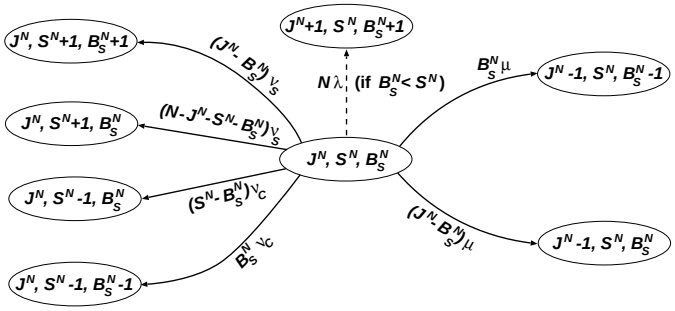

Figure 1: Transitions that bring to a change in B_S for $J \leq N$.

In a similar way, it is possible to show that if $S^N(0)/N$ converges to s_0, $S^N(t)/N$ converges to the solution of the following equation

$$\frac{ds}{dt} = \nu_S - (\nu_S + \nu_C)s(t), \quad s(0) = s_0, \tag{2}$$

and when t diverges $s(t)$ converges to $s^* \triangleq \nu_S/(\nu_S + \nu_C)$, that is the stationary probability that a given datacenter is powered by renewables.

It is clear that we would not have needed fluid models to derive the asymptotic probability that a datacenter is busy or that it is powered by the renewables, but the fluid models allow us to evaluate simply the transient dynamics for the percentage of busy datacenters and of datacenters powered by renewables. Moreover, they are required to characterize the quantity $B_S(t)/N$ that is needed to quantify how many datacenters work using the cheap renewable energy.

In Fig. 1 we show the Markov chain transitions affecting $B_S(t)$, i.e. the number of datacenters working and powered by renewables, when the number of jobs in the system $J(t)$ is smaller than N. As we observed above, for N large enough $J(t) < N$ holds with probability arbitrarily close to one after some finite time depending on $J^N(0)$. For this reason, we can for simplicity assume that the system is in this situation. Observe that the transition indicated in the figure by the dashed line is possible only for specific values of S and B_S. If a new job arrives and there are idle datacenters in state $\mathcal{S}$ (i.e. $B_S < S$) then the job will be assigned to one of them and B_S will increase by one unit. Otherwise B_S will stay constant. If we calculate the drift for $B_S(t)/N$ when $(J(t), S(t), B_S(t)) = (j, s, b_s)$ as done above we obtain that it is equal to

$$g(j, s, b_s) = \begin{cases} \lambda - (\nu_S + \nu_C + \mu)b_s + \nu_S j & \text{if } b_s < s, \\ -(\nu_S + \nu_C + \mu)b_s + \nu_S j & \text{if } b_s = s. \end{cases}$$

Unfortunately the function $g()$ is not continuous, and then neither Lipschitz. Nevertheless, [5] shows that when $B_S^N(0)/N$ converges in probability to $b_{s,0}$,

$B_S^N(t)/N$ is related to the solutions of the following differential inclusion

$$\frac{db_s}{dt} = G(j,s,b_s) = \begin{cases} \{g(j,s,b_s)\} & \text{if } b_s < s, \\ [g(j,s,s), g(j,s,s) + \lambda] & \text{if } b_s = s. \end{cases} \tag{3}$$
$$b_s(0) = b_{s,0}$$

The set-valued function $G(j,s,b_s)$ coincides with $g(j,s,b_s)$ for $b_s < s$, while $G(j,s,s)$ is the interval obtained by the convexification of the accumulation points of $g(j,s,b_s)$ when $b_s = s$. Equation (3) admits at least a solution because $G()$ is upper-semicontinuous and Theorem 5 in [5] shows that in such case

$$\inf_{b_s \in \mathcal{D}} \sup_{0 \le t \le T} \left|\left| \frac{B_S(t)}{N} - b_s(t) \right|\right| \underset{P}{\to} 0,$$

where $\mathcal{D}$ is the set of solutions of Eq. (3). This result has practical utility if we can prove that the differential inclusion (3) has a unique solution. A standard sufficient condition for the uniqueness of the solution is the one side Lipschitz condition [8], that unfortunately does not hold for $G()$. We suspect that Eq. (3) has a unique solution, but we have not been able to prove it. Nevertheless, we can prove that any possible solution converges to the same value as t diverges. This is enough to draw conclusions about the stationary distribution of our stochastic system.

We start observing that for large t $j(t)$ and $s(t)$ are arbitrarily close respectively to the values $j^* = \rho$ and s^*. It holds

$$g(j^*, s^*, s^*) = \nu_S \left(\rho - \frac{\nu_S + \nu_C + \mu}{\nu_S + \nu_C} \right) < 0,$$

because $\rho < 1$. If $\lambda + g(j^*, s^*, s^*) < 0$, then all the values of $G(j^*, s^*, b_s)$ are negative when b_s belongs to an opportune interval $(s^* - \epsilon, s^*]$ and then any possible trajectory of $b_s(t)$ will be constrained to the interval $[0, s^* - \epsilon]$, where the differential inclusion (3) reduces to a usual differential equation with Lipschitz drift and then it admits a unique solution. This solution converges to $(\lambda + \rho\nu_S)/(\nu_S + \nu_C + \mu)$ when t diverges. If $\lambda + g(j^*, s^*, s^*) > 0$, then for $b_s < s^*$

$$g(j^*, s^*, b_s) > \lambda - (\nu_S + \nu_C + \mu)s^* + \nu_S j^* = \lambda + g(j^*, s^*, s^*) > 0,$$

and any trajectory of $b_S(t)$ converges to s^*, that is a stable point because in this case $0 \in G(j^*, s^*, s^*)$. Summarizing, it holds

$$b_s^* \triangleq b_s(\infty) = \begin{cases} s^*, \text{ if } \lambda + g(j^*, s^*, s^*) > 0 \\ (\lambda + \rho\nu_S)/(\nu_S + \nu_C + \mu), \text{ otherwise.} \end{cases}$$

By observing that $\lambda + g(j^*, s^*, s^*) > 0$ is equivalent to $s^* < (\lambda + \rho\nu_S)/(\nu_S + \nu_C + \mu)$, and replacing $\lambda = \rho\mu$ we can write in a more compact way:

$$b_s^* = \min \left\{ s^*, \rho \frac{\nu_S + \mu}{\nu_S + \nu_C + \mu} \right\}. \tag{4}$$

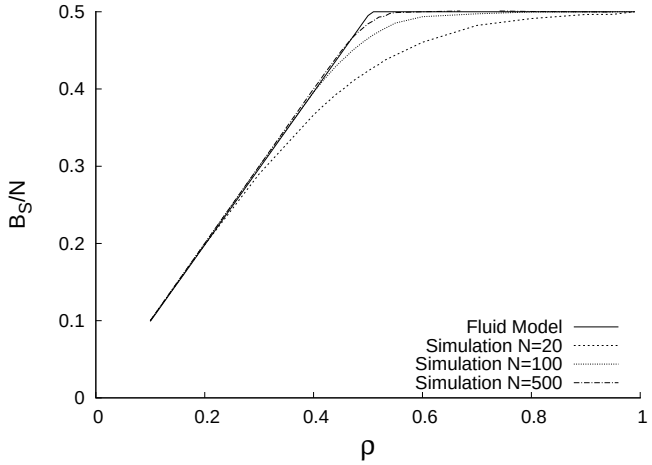

Figure 2: Percentage of datacenters powered by renewables : Fluid Model (b_s^*) vs Simulation Averages ($\nu_S = \nu_C = 0.01$ $\mu = 1$, $N = 20, 100, 500$).

Figure 2 shows how the stationary distribution of B_S^N/N converges to b_s^* as N increases. The quality of the fluid approximation is different for different values of the load ρ. In particular as far as ρ is far from the critical value for which $s^* = \rho(\nu_S + \mu)/(\nu_S + \nu_C + \mu)$, corresponding to the non differentiability in Eq. (4), the approximation is very accurate even for $N = 20$ datacenters. For the critical load, the federation should include an order of magnitude more datacenters to achieve a good level of approximation. For a given value of N the quality of the approximation improves (/worsen) the larger (/smaller) is the acute angle between the two segments determined by the fluid model, as it happens if ν_C increases (/decreases).

5 Exploiting the model

In this section we show how our simple fluid model can help quantifying the potential advantages of a federation of datacenters and the effect of the different parameters.

We start by discussing Eq. (4). The percentage b_s^* of datacenters working and powered by renewables is obviously limited by the percentage s^* of datacenters powered by renewables, and by the percentage ρ of datacenters working, then $b_s^* \leq \min\{s^*, \rho\}$. These two regimes appear also in Eq. (4) and we refer to them as the renewables-limited regime and the load-limited regime. In particular, Eq. (4) shows how close the dispatching algorithm can approach the bound $\min\{s^*, \rho\}$ when, as we assumed, the job will be completed by the datacenter that started working on it. The factor $(\nu_S + \mu)/(\nu_S + \nu_C + \mu)$ multiplying ρ takes into account the fact that a datacenter may change status from $\mathcal{S}$ to $\mathcal{C}$ (or the other way around) after starting to

process a job. These changes limit the utility of job scheduling.

Without the federation every datacenter receives a load ρ and can exploit renewables a fraction s^* of the time. Then the percentage of time a datacenter works and is powered by renewables is ρs^*, that is smaller than b_s^* from Eq. (4):

$$\min\left\{s^*, \rho\frac{\nu_S+\mu}{\nu_S+\nu_C+\mu}\right\} > \rho s^*,$$

because $\rho < 1$ and $(\nu_S+\mu)/(\nu_S+\nu_C+\mu) > \nu_S/(\nu_S+\nu_C) = s^*$. The difference between the left hand side and the right hand side of the inequality times N quantifies how many additional datacenters work powered by renewables thanks to the federation in comparison to the situation when there is no federation. In what follows we compare the corresponding average energy costs, by normalizing the energy cost per time unit to 1 (/0) when the datacenter is (/is not) powered by renewables. The average energy cost per time unit and per datacenter is then:

$$c_f \triangleq \rho - b_s^*, \text{ with the federation}$$
$$c_{nf} \triangleq \rho - \rho s^*, \text{without the federation}$$

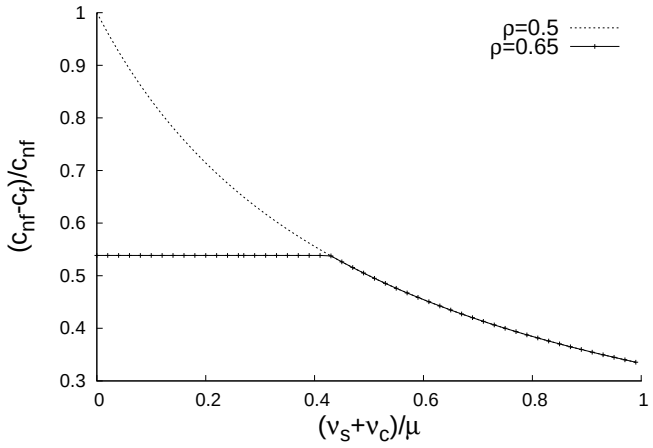

Figure 3: Cost reduction due to the federation vs speed of renewables' dynamics ($s^* = 0.5$).

We focus on the relative cost reduction achieved by the federation in comparison to the uncoordinated case, i.e., on $(c_{nf} - c_f)/c_{nf}$. Fig. 3 shows how the relative cost changes as renewables' dynamics become faster for two different values of the load ρ. We set $\nu_S = \nu_C$ so that s^* is constant and equal to 0.5. Eq. (4) shows that, when ρ and s^* are constant, b_s^* changes only for the effect of the ratio $(\nu_S + \nu_C)/\mu$. In other words, it is not important how fast the quantity of renewable energy produced changes, but how much

faster it changes than the job completion time. Intuitively, if this ratio is very large, the scheduling is not effective, because a datacenter changes its status $\mathcal{S}/\mathcal{C}$ many times before completing the job, so that the job takes advantage of renewables' energy on average a fraction s^* of the time, independently from the status of the datacenter when the job execution started. Fig. 3 shows indeed that the advantage of the federation converges to 0 as the ratio $(\nu_S + \nu_C)/\mu$ diverges. This behaviour is common to both the load values considered. When $\rho = 0.5$, the system is always in the load-limited regime ($b_s^* = \rho(\nu_S + \mu)/(\nu_S + \nu_C + \mu)$) and the advantage of the federation always decreases as the ratio $(\nu_S + \nu_C)/\mu$ increases. When $\rho = 0.65$ the system is initially in the renewables-limited regime, so that the relative gain of the federation is limited by the average availability of renewables' energy and the gain is independent on the speed of their dynamics. This situation corresponds to the initial horizontal part of the corresponding curve. As the speed of renewables' dynamics further increases, the scheduling is no more able to effectively follow them and the system enters in the load-limited regime. The relative improvement from the federation in this regime is independent from ρ, so that both curves in Fig. 3 overlap.

Our analysis shows significant reduction of energy costs achievable by the federation of different datacenters, but, until now, we have assumed that the states of the renewables' sources at the different datacenters are independent. This is not true in general. For example, production from PV panels or wind turbines are clearly positively correlated at nearby locations. When energy quantities produced at the datacenters are positively correlated, the improvement from scheduling is reduced. In order to quantify the effect of positive correlation, we consider the following simple model. We assume that the Markov chain determining the state of a renewable source ($\mathcal{S}$ or $\mathcal{C}$) is modulated by an underlying Markov chain that is common to all the different sources. In particular, as a toy-example, we consider a Markov chain with two states $\mathcal{G}$ and $\mathcal{B}$. The transition rates ν_S and ν_C of each renewable source depend now on the particular state of the modulating Markov chain and we denote them $\nu_{S,G}$, $\nu_{C,G}$, $\nu_{S,B}$ and $\nu_{C,B}$. We consider that

$$s_G^* = \frac{\nu_{S,G}}{\nu_{S,G} + \nu_{C,G}} > \frac{\nu_{S,B}}{\nu_{S,B} + \nu_{C,B}} = s_B^*$$

and then states $\mathcal{G}$ and $\mathcal{B}$ correspond respectively to good and bad weather (at least for the purpose of renewable energy production). It is possible to extend simply our previous analysis, if we assume that the dynamics of the modulating Markov chain are much slower than those of the modulated chain and of job execution (i.e., $\max\{\nu_G, \nu_B\} << \min\{\nu_{S,G}, \nu_{C,G}, \nu_{S,B}, \nu_{C,B}, \mu\}$). In such case, the average percentage of datacenters working and powered by datacenters can be obtained through a weighted sum of what would happen

without modulation as follows

$$b_s^* \approx \pi_G b_{s,G}^* + \pi_B b_{s,B}^*,$$

where $\pi_G = \nu_G/(\nu_G + \nu_B)$, $\pi_B = \nu_B/(\nu_G + \nu_B)$ and $b_{s,G}^*$ (resp. $b_{s,B}^*$) is calculated from Eq. (4) replacing the rates ν_S and ν_C by $\nu_{S,G}$ and $\nu_{C,G}$ (resp. $\nu_{S,B}$ and $\nu_{C,B}$). As we anticipated, the modulating Markov chain correlates the state of the renewable sources. We can quantify this effect by using the correlation coefficient η defined as

$$\eta = \frac{(s_G^*)^2\pi_G + (s_B^*)^2\pi_B - (s_G^*\pi_G + s_B^*\pi_B)^2}{(s_G^*\pi_G + s_B^*\pi_B)(1 - s_G^*\pi_G - s_B^*\pi_B)}.$$

As a sanity check, we observe that if $\nu_{S,G} = \nu_{S,B}$ and $\nu_{C,G} = \nu_{C,B}$ (i.e., the modulating Markov Chain has no effect on the renewables' state evolution), then $\eta = 0$. If instead we have that $s_G^* = 1$ and $s_B^* = 0$, then $\eta = 1$, because all the datacenters are in state $\mathcal{S}$ when the modulating Markov chain is in state $\mathcal{G}$ and in state $\mathcal{C}$ when the modulating Markov chain is in state $\mathcal{B}$.

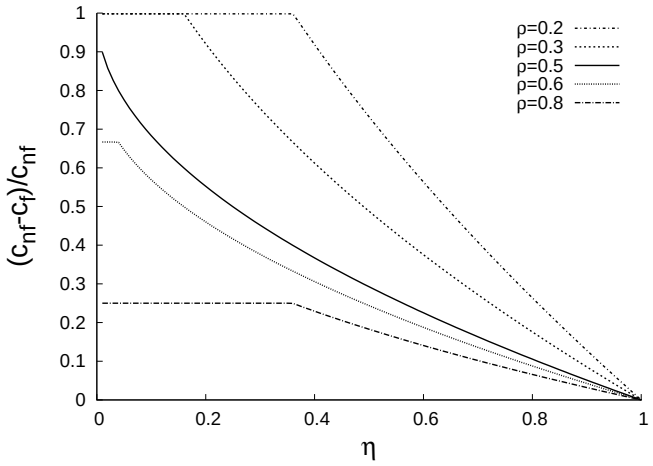

Figure 4: Cost reduction due to the federation vs renewables' correlation ($\nu_G = \nu_B = 0.00001$, $\nu_{S,G} + \nu_{C,G} = \nu_{S,B} + \nu_{C,B} = 0.002$, $\nu_{S,G} = \nu_{C,B}$, $\mu = 1$).

Figure 4 shows the relative cost reduction due to the federation versus the correlation η. In the specific setting considered, the average percentage of time renewables can power datacenters is constant: $s^* = \pi_G s_G^* + \pi_B s_B^* = 0.5$. Then, as the correlation increases s_G^* increases and s_B^* decreases of the same amount. As expected, the benefit from the federation is maximum when renewable sources evolve independently ($\eta = 0$) and null when at any time they are all in the same state ($\eta = 1$). The benefit is non-increasing in η, but, depending on the load ρ, there is a more or less wide range of correlation values for which the benefit does not depend on η. In order to justify this

result, we write the specific expression of b_s^* neglecting for simplicity the rates $\nu_{S,G}$, $\nu_{C,G}$, $\nu_{S,B}$ and $\nu_{C,B}$ when summed to μ, that is much larger. It holds:

$$b_s^* \approx \frac{1}{2}\min\{s_G^*, \rho\} + \frac{1}{2}\min\{s_B^*, \rho\}$$

Under this approximation, the setting $\rho = 0.5 = s^*$ corresponds to the case when the system is at the boundaries between the two regimes for $\eta = 0$. When the correlation increases, the system is i) in the load-limited regime in good weather (state $\mathcal{G}$) with a value $b_{s,G}^*$ almost constant and equal to ρ and ii) in the renewables-limited regime in bad weather (state $\mathcal{B}$) with a value $b_{s,B}^*$ decreasing in η. As a consequence the corresponding curve is decreasing. When $\rho > s^* = 0.5$, the system is the renewables-limited regime in both states $\mathcal{G}$ and $\mathcal{B}$ when $\eta = 0$, and then $b_{s,G}^* = s_G^*$ and $b_{s,B}^* = s_B^*$. As η increases, the increase of s_G^* is exactly compensated by the decrease of s_B^* so that the system exhibits the same relative improvement until η is so large that the system enters in the renewables-limited regime when in bad weather and then the improvement decreases again. Finally, when $\rho < s^* = 0.5$, the system is initially in the load-limited regime in both states, and then $b_{s,G}^* = b_{s,B}^* = \rho$, independent on η. Again, the improvement does not depend on η until η becomes so large that the system enters in the renewables-limited regime when in bad weather.

As we have shown, our simple fluid model reveals the existence of two different regimes and helps to understand and quantify their non-trivial interaction as the parameters change.

6 Conclusions

The paper proposes a model of geographical load balancing strategies for a collection of federated (micro) datacenters powered by renewable energy sources. In our strategy the scheduler uses a selection criterion that prioritizes datacenters where renewable energy is currently produced. For this kind of system we use mean field techniques to derive a simple approximate model that allows us to derive several performance measures. First, asymptotic convergence is proven and the quality of the approximation for finite size systems is evaluated through an ad-hoc simulator. Then, we use the simple fluid model to quantify the effect of the different system parameters and to understand the different tradeoffs.

Acknowledgements

This work was supported in part by the "Investments for the Future" Program reference #ANR-11-LABX-0031-01, funded by the French Government

(National Research Agency, ANR), and in part by the Università Italo-Francese, call Galileo 2015-2016, ref. G15-133.

References

[1] I. J. B. F. Adan, J. Wessels, and W. H. M. Zijm. Analysis of the Asymmetric Shortest Queue Problem. *Queueing Systems*, 8(1):1–58, 1991.

[2] M. Benaïm and J. Y. Le Boudec. A Class of Mean Field Interaction Models for Computer and Communication Systems. *Performance Evaluation*, 65(11-12):823 – 838, 2008.

[3] R. Bianchini. Leveraging Renewable Energy in Data Centers: Present and Future. In *Proc. of HPDC*, 2012.

[4] R. D. Foley and D. R. McDonald. Join the Shortest Queue: Stability and Exact Asymptotics. *The Annals of Applied Probability*, 11(3):569–607, 2001.

[5] N. Gast and B. Gaujal. Markov chains with discontinuous drifts have differential inclusion limits. *Performance Evaluation*, 69(12):623–642, 2012.

[6] G. Guo, Z. Ding, Y. Fang, and D. Wu. Cutting Down the Energy Cost of Geographically Distributed Cloud Data Centers by Using Energy Storage. In *Proc. of GLOBECOM*, 2011.

[7] M. Harchol-Balter, T. Osogami, A. Scheller-Wolf, and A. Wierman. Multi-server queueing systems with multiple priority classes. *Queueing Systems*, 51(3):331–360, 2005.

[8] M. Kunze. *Non-Smooth Dynamical Systems*. Lecture Notes in Mathematics. Springer, 2000.

[9] T. G. Kurtz. Limit Theorems and Diffusion Approximations for Density Dependent Markov Chains. In R. J.-B. Wets, editor, *Stochastic Systems: Modeling, Identification and Optimization, I*, pages 67–78. Springer, 1976.

[10] M. Lin, Z. Liu, A. Wierman, and A. L. H. Online algorithms for geographical load balancing. In *Proc. of IGCC*, 2012.

[11] Z. Liu, M. Lin, A. Wierman, S. Low, and A. L. H. Greening Geographical Load Balancing. In *Proc. of ACM SIGMETRICS*, 2011.

[12] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag, and B. Maggs. Cutting the Electric Bill for Internet-scale Systems. In *Proc. of ACM SIGCOMM*, 2009.

[13] A. Rahman, X. Liu, and F. Kong. A Survey on Geographic Load Balancing Based Data Center Power Management in the Smart Grid Environment. *IEEE Communications Surveys and Tutorials*, 16(1):214–233, 2014.

[14] L. Rao, X. Liu, L. Xie, and W. Liu. Minimizing Electricity Cost: Optimization of Distributed Internet Data Centers in a Multi-Electricity-Market Environment. In *Proc. of INFOCOM*, 2010.

[15] L. Rao, X. Liu, L. Xie, and W. Liu. Coordinated Energy Cost Management of Distributed Internet Data Centers in Smart Grid. *IEEE Transactions on Smart Grid*, 3(1):50–58, March 2012.

[16] Y. Yao, L. Huang, A. Sharma, L. Golubchik, and M. Neely. Data centers power reduction: A two time scale approach for delay tolerant workloads. In *Proc. of INFOCOM*, 2012.

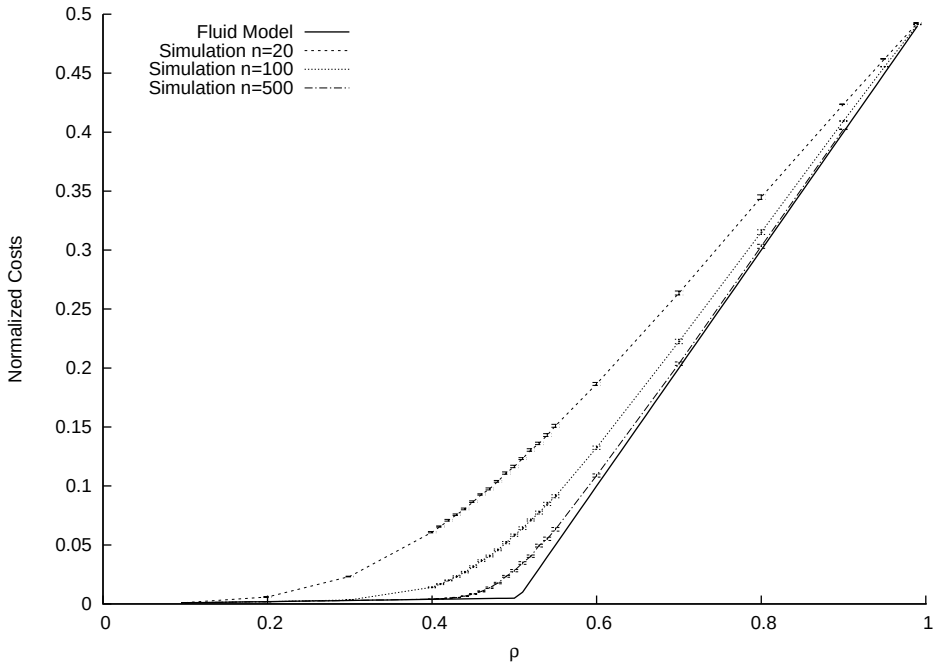
Fluid Model
Simulation n=20
Simulation n=100
Simulation n=500
Normalized Costs
ρ
0
0.05
0.1
0.15
0.2
0.25
0.3
0.35
0.4
0.45
0.5
0
0.2
0.4
0.6
0.8
1

This figure "markov_chain.jpg" is available in "jpg" format from:

http://arxiv.org/ps/1612.03709v1

This figure "markov_chain.png" is available in "png" format from:

http://arxiv.org/ps/1612.03709v1